arXiv:1409.7712v2 [physics.med-ph] 30 Sep 2014

Joint Reconstruction of Multi-channel, Spectral CT Data via Constrained Total Nuclear Variation Minimization

David S. Rigie and Patrick J. La Rivière

October 1, 2014

Abstract

We explore the use of the recently proposed "total nuclear variation" (TNV) [1, 2] as a regularizer for reconstructing multi-channel, spectral CT images. This convex penalty is a natural extension of the total variation (TV) to vector-valued images and has the advantage of encouraging common edge locations and a shared gradient direction among image channels. We show how it can be incorporated into a general, data-constrained reconstruction framework and derive update equations based on the first-order, primal-dual algorithm of Chambolle and Pock. Early simulation studies based on the numerical XCAT phantom indicate that the inter-channel coupling introduced by the TNV leads to better preservation of image features at high levels of regularization, compared to independent, channel-by-channel TV reconstructions.

1 Introduction

Since the influential work of Rudin, Osher, and Fatemi [3], the total-variation (TV) has been used for a variety of image processing tasks, such as denoising [3], deconvolution [4], and demosaicking[5]. Because the TV is a nonsmooth penalty that disfavors oscillations while allowing for sharp discontinuities, it has been found to effectively suppress noise while preserving edges. The TV has more recently been rediscovered as a prior for tomographic inverse problems suffering from incomplete projection data [6, 7]. Consequently, many works have investigated TV-regularized reconstruction algorithms for sparse-view, limited angle, and low-dose CT [8, 9, 10].

In last couple of years, several papers have proposed generalizations of the TV for multi-channel or vector-valued images, in order to extend popular image-processing techniques to the RGB color space [11, 12, 13, 14, 2]. Coincidentally, there is also a fast-growing interest in multi-spectral CT imaging, due to rapid developments in energy-resolving, photon-counting detectors. Since spectral CT is essentially just CT with multiple "color" channels, many of these theoretical advances can be directly applied to spectral CT reconstruction.

The primary consideration when formulating a vectorial TV (VTV) is how the color channels should be coupled. One could naïvely apply the conventional TV to each channel independently, but this would fail to leverage the strong inter-channel correlations inherent to multi-spectral images. In this work, we advocate for a VTV that penalizes the nuclear norm of the Jacobian derivative because it has several desirable properties. This convex prior encourages all image channels to have common edge locations as well as for their gradient

vectors to point in the same direction. The notion that gradient vectors of different color channels should be approximately aligned has been suggested previously [15, 16, 1]. Additionally, we will demonstrate that the inter-channel coupling helps prevent edge smearing in the noisiest color channels, allowing for stronger regularization without significant feature loss. This specific VTV has been recently proposed by Lefkimmiatis [14], Rigie [1], and Holt [2]. We borrow from Holt in referring to it as the "Total Nuclear Variation" (TNV).

The main objective of this paper is to demonstrate the usefulness of the TNV as a regularizer for multi-energy CT reconstruction. Our study contains several important contributions. To our knowledge, this is the first study with empirical results demonstrating the use of the TNV for spectral CT reconstruction, and our results indicate a clear advantage over independent, channel-by-channel reconstructions. Secondly, we propose a novel, noise-balancing transform that improves noise-suppression for multi-channel data for which the color channels have mismatched noise levels. This is a situation that is likely to occur in photon-counting CT, where the lowest and highest energy windows tend to suffer from relatively low count rates. Finally, we demonstrate that the TNV can be applied not only to the energy-binned data, but also to synthesized basis-material data. In most color image processing tasks, the measured RGB basis is the same as the observed basis. The analog in spectral CT would be to directly reconstruct the energy-bin data, so we will refer to this as the "color" basis. However, often the goal is, instead, to reconstruct density maps of different basis materials [17, 18], which can be obtained through a non-linear transformation of the measured data [19]. In this work, we will refer to this space as the "material" basis. We show empirically that the TNV works well in either basis and that it may actually be advantageous to co-reconstruct the color basis and the material basis. We refer to this as the "hybrid" basis, which is formed by concatenating the color and material basis channels. By simultaneously reconstructing the low-noise, color basis and the high-noise, material basis, the TNV is able to better suppress much of the noise that is amplified during the decomposition.

First, we will review some theoretical concepts and define the total nuclear variation (TNV). Next, we will describe our data-constrained, TNV reconstruction model and show how the first-order, primal-dual algorithm of Chambolle and Pock [20] can be used to derive efficient update equations. We compare our TNV reconstruction model to performing separate channel-by-channel reconstructions using simulated, dual-energy and photon-counting data using the anthropomorphic XCAT phantom [21]. In the photon-counting study, we present reconstruction results from the color basis, where the color channels correspond to five equispaced energy windows. In the dual kVp study, we present reconstructions of the hybrid basis, where 80 kVp, 140 kVp, bone, and soft-tissue images are simultaneously reconstructed. Overall, we find that the TNV regularization results in less edge smearing at the same noise levels, while the computation time is very similar to the naïve channel-by-channel reconstruction. Additionally, the hybrid TNV reconstruction technique seems to be a promising method for controlling noise in the basis-material images.

2 Theory

2.1 Notation and definitions

In this work, we deal only with 2D images, but the generalization to 3D is straightforward. Consider a single-energy CT image vector $u \in \mathbb{R}^{\mathcal{L}}$, where $\mathcal{L}$ is the total number of image pixels. We will also refer to its discrete derivative $Du \in \mathbb{R}^{2\cdot\mathcal{L}}$, which consists of a 2D derivative vector $[Du]_\ell$ at every pixel location. where $[.]_\ell$ denotes the a particular pixel-index. This quantity is also directly related to the discrete gradient by the

following relationship:

$$[Du]_\ell^T = [\nabla u]_\ell \,. \tag{1}$$

The quantities u and Du are both 1D vectors, so we will use $\|.\|$ and $\langle ., . \rangle$ to refer to the usual vector norm and inner product. We will also refer to a multi-channel image, $\boldsymbol{u} \in \mathbb{R}^{\mathcal{M} \cdot \mathcal{L}}$, with $\mathcal{M}$ spectral channels and its discrete derivative $D\boldsymbol{u} \in \mathbb{R}^{(\mathcal{M} \times 2) \cdot \mathcal{L}}$. The quantity $D\boldsymbol{u}$ is a discrete analog to the Jacobian derivative of a vector-function, and at every pixel location the first-order derivatives are fully characterized by the $\mathcal{M} \times 2$ rectangular matrix $[D\boldsymbol{u}]_\ell$. For two tensors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{(\mathcal{M} \times 2) \cdot \mathcal{L}}$, we define the inner product $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \langle \operatorname{vec} \boldsymbol{x}, \operatorname{vec} \boldsymbol{y} \rangle$, where vec(.) unfolds a tensor into a 1D vector.

2.2 The scalar total variation

Total variation regularized CT reconstruction algorithms have been studied extensively due to the approximate gradient sparsity of CT images. Many works have demonstrated that such schemes can often yield high quality images from severely undersampled projection data [6, 7, 22, 8, 23].

The anisotropic TV The "anisotropic" TV is simply defined as the ℓ_1-norm of the derivative of the image u,

$$\mathrm{TV}_a(u) = \|Du\|_1, \tag{2}$$

where it is useful to think of the $\|.\|_1$ as a "sparsity-inducing" norm because it is a convex relaxation of $\|.\|_0$.

The isotropic TV Somewhat more common is the "isotropic" TV, which is defined in terms of the gradient-magnitude image,

$$\mathrm{TV}(u) = \sum_{\ell=1}^{\mathcal{L}} \| \, [Du]_\ell \, \|_2 \tag{3}$$

This penalty function has similar properties to the anisotropic TV, but is rotationally invariant, whereas the anisotropic TV tends to under-penalize horizontal and vertical edges.

2.3 The vectorial total variation

Now, we consider generalizing the TV to the $\mathcal{M}$-channel image $\boldsymbol{u}$.

Naïve generalization to multi-channel images The simplest way to extend the definition of the TV to this vector-valued image, $\boldsymbol{u}$, would be to take the summation of the total-variation of each channel. This is given by

$$\mathrm{TV}_S(\boldsymbol{u}) = \sum_{m=1}^{\mathcal{M}} \mathrm{TV}(\boldsymbol{u}_m). \tag{4}$$

Since this penalty is additively separable, this would be akin to just reconstructing each channel separately, so we will refer to this approach as the "channel-by-channel" TV.

A class of vectorial TV regularizers. Now, we will describe a more principled way to generalize the TV to multi-channel images. We mentioned earlier that the first-order derivatives of $\boldsymbol{u}$ can be characterized

at every pixel-location by the Jacobian matrix $[D\boldsymbol{u}]_\ell$. An interesting class of regularizers is described by

$$\mathrm{VTV}(\boldsymbol{u}) = \sum_{\ell=1}^{\mathcal{L}} \| \, [D\boldsymbol{u}]_\ell \, \|_{S_p}, \tag{5}$$

where the matrix norm $\|.\|_{S_p}$ is known as the "Schatten p-norm," which is equivalent to the ℓ_p-norm of the singular values, $\|X\|_{S_p} = \|\sigma(X)\|_p$, where $\sigma(X)$ is a vector comprised of the singular values of matrix X. Some special cases include S_∞ (spectral norm), S_2 (Frobenius norm), and S_1 (nuclear norm). All three of these norms have been suggested as potential VTV penalties [12, 13, 14, 2, 1], and they share a few important properties. Because they are norms defined by the singular values of the Jacobian, they are rotationally invariant and convex. Additionally, for a single-channel image, they all reduce exactly to the conventional, isotropic TV.

The Total Nuclear Variation (TNV) The total nuclear variation (TNV), defined as the nuclear norm of the Jacobian, is the only one of these VTV's that encourages the gradient vectors of all of the image channels to point in a common direction. We have found empirically that this is a very useful prior for spectral CT images, and that it generally outperforms the other VTV penalties. Therefore, in this work we will focus solely on the TNV, which is defined by

$$\mathrm{TNV}(\boldsymbol{u}) = \sum_{\ell=1}^{\mathcal{L}} \| \, [D\boldsymbol{u}]_\ell \, \|_{S_1}. \tag{6}$$

To see how this TNV couples the different energy channels of $\boldsymbol{u}$, note that the Jacobian at pixel ℓ can also be expressed in terms of the gradient vectors of the various image channels:

$$[D\boldsymbol{u}]_\ell = \begin{bmatrix} \longleftarrow & [\nabla \boldsymbol{u}_1]_\ell^T & \longrightarrow \\ \longleftarrow & \vdots & \longrightarrow \\ \longleftarrow & [\nabla \boldsymbol{u}_\mathcal{M}]_\ell^T & \longrightarrow \end{bmatrix}_{\mathcal{M}\times\mathcal{N}}. \tag{7}$$

The nuclear norm encourages sparsity in the singular values of $[D\boldsymbol{u}]_\ell$ [24, 25]. If pixel ℓ lies within a

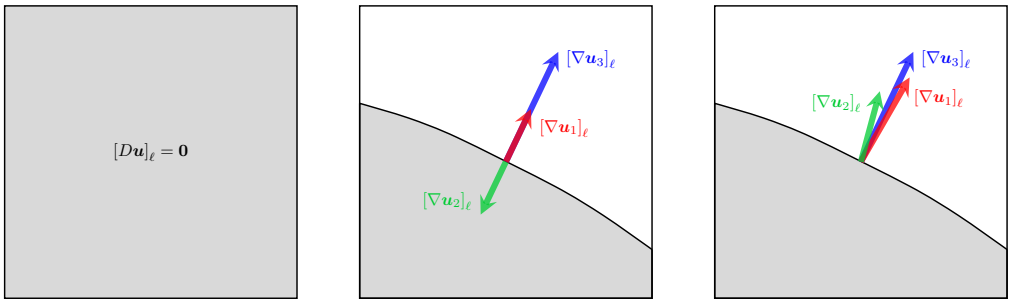

(a) All singular values are zero

(b) 1 non-zero singular value

(c) All non-zero singular values

Figure 1: This illustrates the relationship between the directionality of the gradient vectors of a 3-channel image and the singular values of its Jacobian. When all image channels are constant valued, all singular values of the Jacobian are zero (a). If the gradient vectors are parallel or anti-parallel, one singular value is non-zero (b). All singular values will be nonzero if the gradients have unique directions (c).

constant-valued region of all of the image channels then this Jacobian matrix will be entirely null-valued, and all of its singular values will be zero. Further, if all of the gradient vectors of the various image channels are parallel or anti-parallel, such as at a common edge, then the Jacobian will be rank one and thus have only one non-zero singular value. This is because parallel or anti-parallel vectors are not linearly independent. The fact that the TNV treats parallel and anti-parallel vectors the same means that it is robust to contrast inversions. We also point out that if one of the gradient vectors has zero magnitude, the rank is still one, so the TNV also appears to be robust in handling edges that may not exist in every image channel. In sum, encouraging singular-value sparsity in the Jacobian is akin to encouraging common edge locations and gradient directions. This notion is illustrated graphically in Figure 1.

Converting to a "noise-balanced" image space In spectral CT imaging it is often the case that certain color channels are significantly noisier than others, and we have found the TNV is much more effective if a noise-balancing transform is applied prior to reconstruction or denoising. For example, suppose we had a two-channel image, $\boldsymbol{u}$, with zero-mean, Gaussian noise with variances given by σ_1^2 and σ_2^2 for channels 1 and 2, respectively. If we wanted to use the TNV to denoise $\boldsymbol{u}$, we should first define a noise-balanced image $\boldsymbol{u'}$ as in equation 8.

$$\boldsymbol{u'} = \begin{bmatrix} u_1/\sigma_1 \\ u_2/\sigma_2 \end{bmatrix} \tag{8}$$

Each channel is globally scaled so that the per-channel noise is the same. After denoising, the inverse scale factors are applied in order to return to the original image space. For image reconstruction, the same type of scaling would be applied to the sinogram data. In general, the projection data will not have a uniform noise level within a single color channel, so we use some average measure of channel noise to determine the global scale factor. We find that this noise-balancing procedure is a very important step, as it greatly improves the noise suppression in multi-spectral images with unequal noise levels. In many cases, even images that are corrupted by extremely high noise levels may be well recovered using this technique as long as at least one of the color channels is reasonably clean.

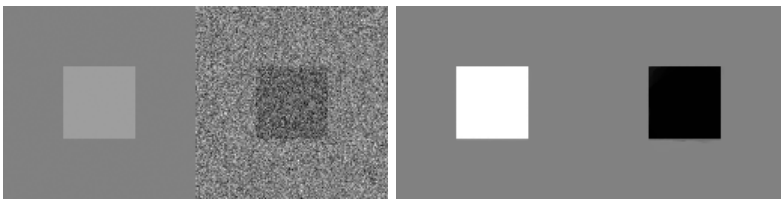

Figure 2: A noisy two-channel image (left) with grossly mismatched noise levels. The image is successfully denoised using the TNV in conjunction with the noise-balancing procedure (right). Even the extremely noisy channel is recovered nearly perfectly due to the tight coupling to the clean image channel.

3 Materials and methods

3.1 General reconstruction model

Assume we have a set of $\mathcal{M}$ sinograms, corresponding either to the directly measured energy channels of a spectral CT system (color basis) or to the material basis projections. We jointly reconstruct these channels

by solving the following data-constrained optimization problem:

$$\begin{aligned} \min_{\boldsymbol{u}} \quad & \mathrm{VTV}(\boldsymbol{u}) \\ \text{s.t.} \quad & \|A\boldsymbol{u} - \boldsymbol{f}\|_W \leq \epsilon, \end{aligned} \tag{9}$$

where $\boldsymbol{f}$ is the multi-channel projection data related to $\boldsymbol{u}$ by the discrete fan-beam or cone-beam projection operator A. The projector is defined on a per channel basis,

$$(A\boldsymbol{u})_m = A_m \boldsymbol{u}_m, \tag{10}$$

where A_m characterizes the measurement geometry associated with channel m. The norm on the data constraint $\|.\|_W$ is a weighted ℓ_2 norm, defined by

$$\|x\|_W^2 = \langle x, Wx \rangle . \tag{11}$$

This penalized weighted least squares (PWLS) data model is often used in CT reconstruction because when the weights are selected such that $W = K_{\boldsymbol{f}}^{-1}$, the data fidelity term is a maximum likelihood estimator for Gaussian data with covariance $K_{\boldsymbol{f}}$. The single adjustable parameter ϵ controls the balance between data-fidelity and regularity. By fixing ϵ we can directly compare reconstructions using the naïve channel-by-channel TV_S to our proposed TNV subject to the same data fidelity constraint.

3.2 The Chambolle Pock Primal Dual Algorithm

In this section, we will provide an overview of the first-order, primal-dual algorithm of Chambolle and Pock [20] and demonstrate how to apply it to our reconstruction model. For a tutorial on how to apply the CPPD algorithm to various CT reconstruction schemes, we refer the reader to Sidky et al [26].

The proximal mapping In order to describe the CPPD algorithm, we must first introduce the concept of a proximal mapping. For a function $f(x)$, the proximal mapping is defined by

$$\mathbf{prox}_f(x) = \operatorname*{argmin}_u \; f(u) + \frac{1}{2}\|u - x\|_2^2. \tag{12}$$

It will also be useful to consider the proximal mapping of the scaled function λf, which can be written as

$$\mathbf{prox}_{\lambda f}(x) = \operatorname*{argmin}_u \; f(u) + \frac{1}{2\lambda}\|u - x\|_2^2. \tag{13}$$

The proximal mapping can be interpreted as a generalized projection operator because for the special case that $f(x)$ is an indicator function, then it is precisely a euclidean projection. For an extensive overview of the prox operator and applications, refer to Parikh and Boyd [27].

A general saddle-point problem Let X and Y be two finite-dimensional, real vector spaces. The CPPD algorithm is designed to solve the saddle-point problem described by

$$\min_{x \in X} \max_{y \in Y} \; \langle Kx, y \rangle - F(y) + G(x), \tag{14}$$

where $F(y)$ and $G(x)$ are convex functions with very simple proximal mappings, and K is a general linear operator. In particular, the proximal mappings associated with F and G should have a closed form or be easily solvable to high precision. It turns out that many interesting image processing and reconstruction problems fit this description, including the VTV reconstruction model already described in the previous section. The CPPD update equations are summarized by Algorithm 1. The parameters θ, σ, τ can be

Algorithm 1 CPPD Algorithm

Initialize: $\theta \in [-1, 1]$, $\sigma\tau\|K\|_2^2 < 1$, $x^{(0)}, \bar{x}^{(0)} \in X$, $y^{(0)} \in Y$ ▷ $\|K\|_2 \equiv \sigma_{\max}(K)$
repeat
  $y^{(k+1)} = \mathbf{prox}_{\sigma F}\left(y^{(k)} + \sigma K \bar{x}^{(k)}\right)$
  $x^{(k+1)} = \mathbf{prox}_{\tau G}\left(x^{(k)} - \tau K^T y^{(k+1)}\right)$
  $\bar{x}^{(k+1)} = x^{(k+1)} + \theta\left(x^{(k+1)} - x^{(k)}\right)$
until $x^{(k+1)}, y^{(k+1)} = x^{(k)}, y^{(k)}$ ▷ stop when convergence criteria met

thought of as step-size parameters that affect convergence speed but not the final solution, and $\|K\|_2$ is the "spectral norm" of the operator K, which is equivalent to its largest singular value. The power method, outlined in Sidky et al [26], provides a very efficient iterative procedure for computing $\|K\|_2$ for very large-scale problems, involving only matrix-vector multiplications. This algorithm has close ties with several other popular algorithms, including ADMM, split-bregman, and proximal-point. This is further discussed in Chambolle and Pock [20].

Applying the CPPD algorithm to VTV reconstruction Now we will outline how the CPPD algorithm can be applied to the data-constrained optimization problem of (9). First we rewrite (9) as

$$\min_{\boldsymbol{u}} \quad \mathrm{VTV}(\boldsymbol{u}) + \delta_\epsilon\left(A\boldsymbol{u} - \boldsymbol{f}\right), \tag{15}$$

where the indicator function $\delta_\epsilon(\boldsymbol{x})$ is defined by

$$\delta_\epsilon(\boldsymbol{x}) = \begin{cases} 0 & \|\boldsymbol{x}\|_W \leq \epsilon \\ \infty & \|\boldsymbol{x}\|_W > \epsilon \end{cases}. \tag{16}$$

Now, we introduce an auxiliary variable $\boldsymbol{q}$ and rewrite the data constraint as an optimization problem:

$$\delta_\epsilon(A\boldsymbol{u} - \boldsymbol{f}) = \max_{\boldsymbol{q}} \ \langle \boldsymbol{q}, A\boldsymbol{u} - \boldsymbol{f}\rangle - \epsilon\|W^{-1/2}\boldsymbol{q}\|_2. \tag{17}$$

Similarly, we can rewrite the VTV penalty function as an optimization over another auxiliary variable, $\boldsymbol{z}$. For the channel-by-channel TV, we have

$$\mathrm{TV}_S(\boldsymbol{u}) = \max_{\boldsymbol{z}} \ \langle D\boldsymbol{u}, \boldsymbol{z}\rangle - \delta_{\mathcal{S}}(\boldsymbol{z}), \tag{18}$$

where the definition of the set $\mathcal{S}$ is given by

$$\mathcal{S} = \left\{ \boldsymbol{z} \in \mathbb{R}^{(\mathcal{M}\times 2)\cdot\mathcal{L}} \ : \ \|\left[\boldsymbol{z}_m\right]_\ell\|_2 \leq 1 \ \forall \ m \in 1, \ldots, \mathcal{M}, \ ; \ell \in 1, \ldots, \mathcal{L} \right\}. \tag{19}$$

The quantity $[\boldsymbol{z}_m]_\ell$ refers to image channel m at pixel ℓ. Similarly, we can write the proposed TNV penalty function in terms of an optimization over auxiliary variable $\boldsymbol{z}$:

$$\text{TNV}(\boldsymbol{u}) = \max_{\boldsymbol{z}} \ \langle D\boldsymbol{u}, \boldsymbol{z}\rangle - \delta_N(\boldsymbol{z}), \tag{20}$$

where the set N is defined by

$$N = \left\{ \boldsymbol{z} \in \mathbb{R}^{(\mathcal{M}\times 2)\cdot\mathcal{L}} \ : \ \sigma_{\max}([\boldsymbol{z}]_\ell) \leq 1 \ \forall \ \ell \in 1, \ldots, \mathcal{L} \right\}. \tag{21}$$

The quantity $[\boldsymbol{z}]_\ell$ is an $\mathcal{M} \times 2$ matrix, with the same dimensions as the Jacobian derivative at pixel ℓ, and $\sigma_{\max}$ is its largest singular value. These transformations, which are detailed in Appendix A, allow us to write our original VTV reconstruction model as a primal-dual, saddle-point problem,

$$\min_{\boldsymbol{u}} \max_{\boldsymbol{z},\boldsymbol{q}} \ \langle D\boldsymbol{u}, \boldsymbol{z}\rangle + \langle \boldsymbol{q}, A\boldsymbol{u} - \boldsymbol{f}\rangle - \delta_{\mathcal{S}/N}(\boldsymbol{z}) - \epsilon\|W^{-1/2}\boldsymbol{q}\|_2, \tag{22}$$

where the indicator function is either $\delta_{\mathcal{S}}$ for the channel-by-channel TV_S or δ_N for the proposed TNV. We can directly apply the CPPD update equations of Algorithm 1 by making the following assignments.

$$\begin{gathered} x \to \boldsymbol{u}, \quad y \to \begin{pmatrix} \boldsymbol{q} \\ \boldsymbol{z} \end{pmatrix}, \quad K \to \begin{pmatrix} A \\ D \end{pmatrix} \\ F(\boldsymbol{q}, \boldsymbol{z}) = \delta_{\mathcal{S}/N}(\boldsymbol{z}) - \epsilon\|W^{-1/2}\boldsymbol{q}\|_2 - \langle \boldsymbol{q}, \boldsymbol{f}\rangle, \quad G(\boldsymbol{u}) = 0 \end{gathered} \tag{23}$$

The resulting update equations are given by Algorithm 2. The operator $\Pi_{\mathcal{S}/N}$ is a Euclidean projection

Algorithm 2 Data-constrained VTV Update Equations

1: Initialize: $\theta = 1, \ \sigma\tau \left\| \begin{pmatrix} A \\ D \end{pmatrix} \right\|_2^2 = 1, \ (\boldsymbol{u}, \bar{\boldsymbol{u}}, \boldsymbol{q}, \boldsymbol{z}) \to \boldsymbol{0}$
2: **repeat**
3: $\quad \boldsymbol{z}^{(k+1)} = \Pi_{\mathcal{S}/N}\left(\boldsymbol{z}^{(k)} + \sigma D\bar{\boldsymbol{u}}^{(k)}\right)$
4: $\quad \boldsymbol{q}^{(k+1)} = \mathbf{prox}_{\sigma\epsilon F_q}\left(\boldsymbol{q}^{(k)} + \sigma\left(A\bar{\boldsymbol{u}}^{(k)} - \boldsymbol{f}\right)\right)$
5: $\quad \boldsymbol{u}^{(k+1)} = \boldsymbol{u}^{(k)} + \tau\left(\text{Div}\, \boldsymbol{z}^{(k+1)} - A^T\boldsymbol{q}^{(k+1)}\right)$
6: $\quad \bar{\boldsymbol{u}}^{(k+1)} = \boldsymbol{u}^{(k+1)} + \theta\left(\boldsymbol{u}^{(k+1)} - \boldsymbol{u}^{(k)}\right)$
7: **until** ▷ stop when convergence criteria met

onto the set $\mathcal{S}$ or N, and the Div operator is the negative transpose of the discrete derivative operator D. Note that the $\boldsymbol{q}$-update is written in terms of the proximal map of the function F_q, which we define as $F_q = \|W^{(-1/2)}\boldsymbol{q}\|_2$. While there does not appear to be a closed-form expression for this proximal map, it can be reduced to a very simple scalar optimization problem. This is detailed in the appendix along with an efficient method for performing the Euclidean projection in the primal variable update.

3.3 Simulation studies

To investigate the impact of the proposed vectorial TV regularization, we performed two numerical simulation studies with the computerized XCAT phantom. In the first experiment we simulate an ideal, photon-counting system with 5 energy windows and directly reconstruct images corresponding to the log-normalized bin data. We will refer to this image basis as the "color" basis because each image channel corresponds directly to the

measurements of one energy window.

In the second experiment, we simulate an ideal, dual-kVp scan and perform what we will refer to as a "hybrid" reconstruction. As we will detail later, we first decompose the 80/140 kVp data into a bone/soft-tissue "material" basis and then we co-reconstruct this synthetic data with the raw 80/140 kVp sinograms. The TNV penalty couples all four image channels so that the relatively noisy basis-material channels may benefit from the higher SNR color channels.

3.3.1 Data generation

The XCAT shoulder phantom All simulations used the same 2D pixelized shoulder phantom, which was generated from an axial slice of the XCAT phantom and is pictured in Figure 3. The XCAT software package was used to generated a set of bone and soft-tissue density maps on a 2048×2048 pixel grid. These density maps were then used as input to a polyenergetic, distance-driven projector model to generate the simulated raw data that would be inputted to our reconstruction algorithm.

Ideal photon counting model To generate the 5-bin photon-counting data, we used a realistic 120 kVp simulated x-ray tube spectrum with hard thresholds at 40,60,80,100, and 120 keV. We did not model any physical factors in the detector response, so our bin response functions are perfect rect functions. Our simulated spectrum had virtually no emissions below 20 keV, so we can think of these bins as being evenly spaced. Using a distance-driven projector model, polyenergetic projections were computed with 896 detector elements per view and 400 views, and Poisson noise was added.

Dual kVp hybrid model The other configuration we looked at was an 80/140 dual kVp acquisition with the same number of detectors and views as the photon counting simulation. We simulated realistic 80 and 140 kVp x-ray tube spectra to generate two sets of consistent projection data. This type of consistent, dual-kVp data can be acquired on many current scanners by performing two back-to-back scans or can be estimated from a fast kV-switching geometry. Gaussian noise was added to approximate a compound Poisson noise model. From the noisy 80 and 140 kVp projection data, we synthesized bone and soft-tissue basis sinograms using a maximum-likelihood material decomposition. In the "hybrid" study, we will co-reconstruct these 4 channels of data, consisting of our log-normalized, dual kVp data (color basis) and the synthesized, material-basis data.

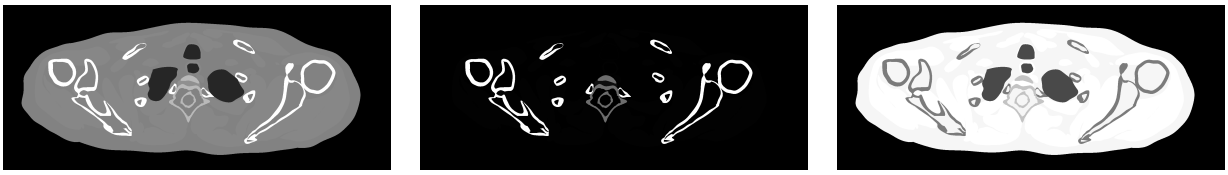

(a) 60 keV reference image (b) bone density map (c) soft-tissue density map

Figure 3: The 2D XCAT shoulder phantom used for all simulations studies, depicted at 60 keV (a) for reference. The projection data were generated using bone (b) and soft-tissue (c) density maps with the appropriate mass-attenuation curves.

3.3.2 Preprocessing and reconstruction

Both the 5-channel photon-counting data and the 4-channel dual kVp hybrid data were reconstructed using the reconstruction model outlined in (9). For reconstruction, we used a 512×512 grid of 1 mm pixels and a matched projector/backprojector pair based on Joseph's method.

Computing data-weights For the raw-basis projection data, we will use the data-weighting approach described by Bouman [28], which results in a diagonal W matrix, where the diagonal elements equal the pre-logged projection data. This is a quadratic approximation to the log-likelihood. For the synthetic, bone/soft-tissue sinograms, we estimate the variances using the Cramer-Rao lower bound [29] and weight by their inverse. This is similar to the approach described by Schirra [30] and Sawatzky [31], but we ignore the off-diagonal terms in this work for simplicity. We found that having a non-diagonal weighting matrix made it immensely more complicated to derive the algorithm update equations when using this data-constrained form. However, in practice, it is often easier to solve the regularized (unconstrained) problem, and incorporating the off-diagonal covariance terms poses no challenge in this case. We focus on the data-constrained form in this work, simply because it provides a straightforward way to compare the vectorial TV to the channel-by-channel TV.

Tuning the data-constraint parameter The only parameter that we will vary in our reconstruction model (9) is ϵ, which controls the trade-off between data fidelity and smoothness. In general, smaller values of ϵ will result in noisier images, while higher values allow the VTV term to find a smoother solution. At the extreme case of $\epsilon = 0$, only images with projections that exactly match the measured data are allowed. There is also some finite value $\epsilon = \epsilon_{\max}$ for which the algorithm will return an image of all zeros. In order to determine an interesting range for ϵ, we first compute a ground truth image $\boldsymbol{u}_{\text{true}}$ by performing FBP on noiseless, non-sparse projection data. Then we define a new parameter ϵ^* which is defined by equation 24,

$$\epsilon^* \equiv \|A\boldsymbol{u}_{\text{true}} - \boldsymbol{f}\|_W \tag{24}$$

where $\boldsymbol{f}$ are the noisy projection data we wish to reconstruct. In this study we will select values of ϵ that satisfy $\epsilon = \alpha\epsilon^*$, where $\alpha \in (0, 1)$. We subjectively selected a range of α values that represent a range of solutions from under-smoothed to over-smoothed in order to demonstrate how the TNV compares to the channel-by-channel TV in various regimes. As described earlier, the projection data were re-scaled prior to reconstruction, so that the average noise levels were approximately the same in every spectral channel.

4 Results

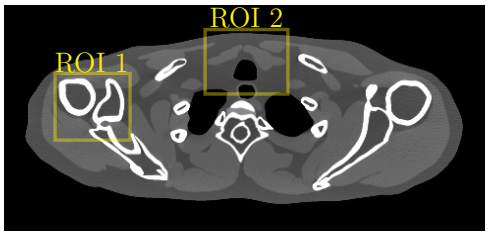

Figure 4: XCAT reference image, with ROI's indicated in yellow

4.1 Photon counting study

Here we present the results of performing a simultaneous reconstruction of the five photon-counting bins after log-normalization, using both the channel-by-channel TV_S and the proposed TNV. This is an example of performing a joint reconstruction on data in the color basis because we did not perform a material decomposition. In this setup, bin 1 (0-40 keV) was the noisiest, due to significant attenuation below 40 keV. We find that the inter-channel coupling introduced by the total nuclear variation has the greatest impact on the noisiest channels, so we present reconstructed images from this energy bin, focusing on the ROI's indicated in Figure 4. The resulting images are depicted in figures 5 through 6. The image window was manually adjusted for each ROI to highlight the relevant structures but is fixed within a particular figure. In general, we find that the images reconstructed with TNV regularization suffer from less edge blurring as the smoothing parameter ϵ is increased. The profile plot in Figure 7 gives a closer look at how the TNV better preserves bony structures in the lowest energy bin image. We confirmed that these profiles were extracted from images of similar noise levels by measuring the sample variance in a nearby, uniform muscle ROI. This is expected because our data-constrained, reconstruction model allows us to select the noise level directly by tuning the ϵ parameter.

(a) TV_S $\epsilon = 0.001\epsilon^*$ RMSE= 0.0444 | (b) TV_S, $\epsilon = 0.0013\epsilon^*$ RMSE= 0.0436 | (c) TV_S, $\epsilon = 0.0016\epsilon^*$ RMSE= 0.0450 | (d) TV_S, $\epsilon = 0.002\epsilon^*$ RMSE= 0.0484 | (e) TV_S, $\epsilon = 0.003\epsilon^*$ RMSE= 0.0609

(f) TNV, $\epsilon = 0.001\epsilon^*$ RMSE= 0.0390 | (g) TNV, $\epsilon = 0.0013\epsilon^*$ RMSE= 0.0347 | (h) TNV, $\epsilon = 0.0016\epsilon^*$ RMSE= 0.0332 | (i) TNV, $\epsilon = 0.002\epsilon^*$ RMSE= 0.0327 | (j) TNV, $\epsilon = 0.003\epsilon^*$ RMSE= 0.0456

Figure 5: Bin 1 (0-40keV) image, bone ROI/window comparison with channel-by-channel TV (top) and TNV (bottom). Grayscale window in cm^{-1} $[0.30, 0.85]$.

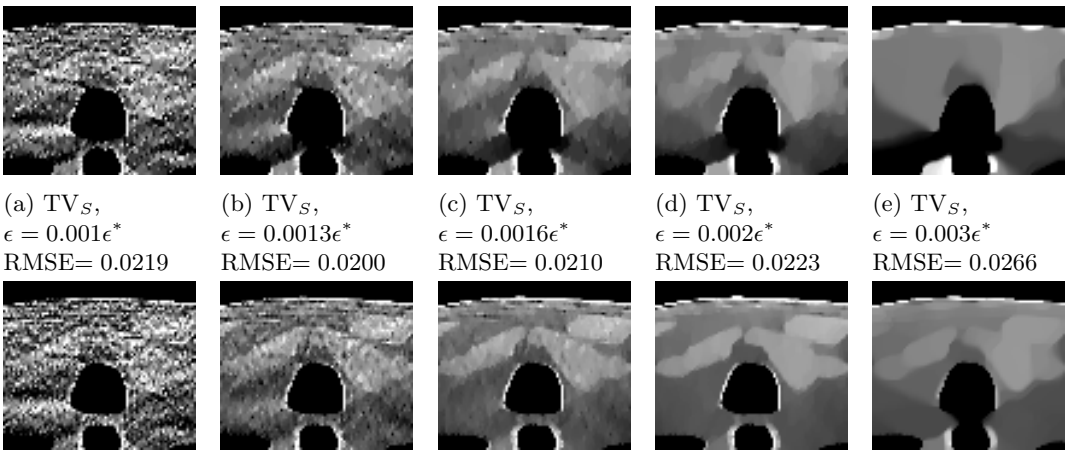

(a) TV_S, $\epsilon = 0.001\epsilon^*$ RMSE= 0.0219
(b) TV_S, $\epsilon = 0.0013\epsilon^*$ RMSE= 0.0200
(c) TV_S, $\epsilon = 0.0016\epsilon^*$ RMSE= 0.0210
(d) TV_S, $\epsilon = 0.002\epsilon^*$ RMSE= 0.0223
(e) TV_S, $\epsilon = 0.003\epsilon^*$ RMSE= 0.0266

(f) TNV, $\epsilon = 0.001\epsilon^*$ RMSE= 0.0216
(g) TNV, $\epsilon = 0.0013\epsilon^*$ RMSE= 0.0176
(h) TNV, $\epsilon = 0.0016\epsilon^*$ RMSE= 0.0181
(i) TNV, $\epsilon = 0.002\epsilon^*$ RMSE= 0.0189
(j) TNV, $\epsilon = 0.003\epsilon^*$ RMSE= 0.0175

Figure 6: Bin 1 (0-40keV) image, soft-tissue ROI comparison with channel-by-channel TV (top) and TNV (bottom). Grayscale window in cm^{-1} $[0.30, 0.35]$.

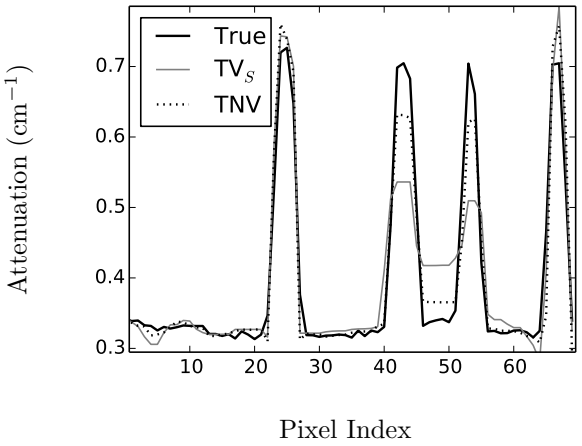

Figure 7: This is a vertical profile through the 0-40 keV, bone-ROI image with $\epsilon = 0.0016\epsilon^*$. The TNV regularized reconstruction shows slightly better preservation of bony structures. The noise levels were estimated from a nearby ROI in a uniform muscle region ($\sigma_{TV_S} = 0.00245$, $\sigma_{\mathrm{TNV}} = 0.00221$).

4.2 Dual kVp Hybrid Study

In this study, we co-reconstruct the synthetic soft-tissue/bone image channels with the raw 80/140 kVp data. In general, the material decomposition greatly amplifies noise, so that the soft-tissue and bone images tend to have a much lower CNR than the raw 80 and 140 kVp images. This noise amplification is due to the ill-conditioning of the inversion step in the basis change, which is caused by the relatively poor spectral separation between basis materials. This poor spectral separation also explains the strong negative correlation between the synthesized material channels. We hypothesize that by coupling the high CNR raw image channels and the low CNR synthetic image channels, we may be able to improve noise suppression in the synthetic data. We call this technique of reconstructing the synthetic and raw data simultaneously "hybrid" reconstruction. We present ROI's from both the bone density image and the soft-tissue density image over a range of different ϵ values in Figures 8 and 9. We find that using the TNV allows for a high

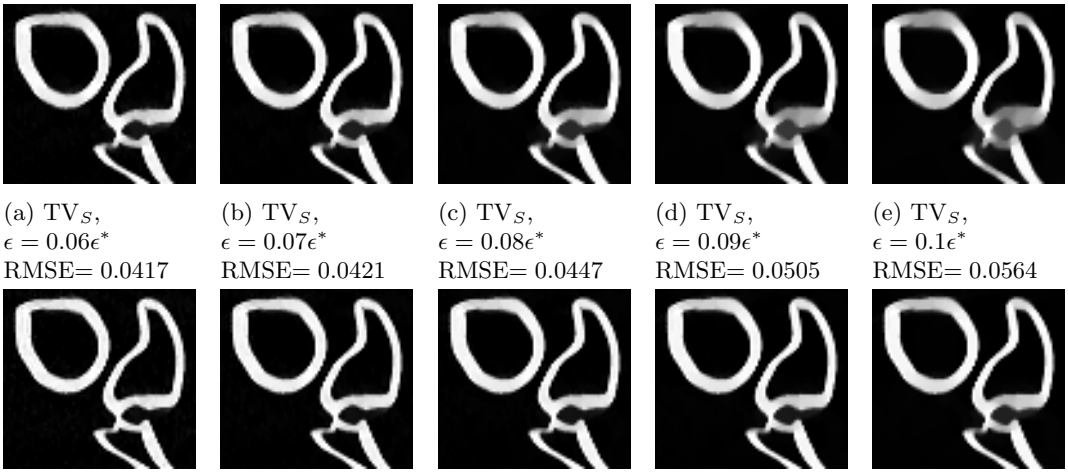

(a) TV_S, $\epsilon = 0.06\epsilon^*$ RMSE= 0.0417
(b) TV_S, $\epsilon = 0.07\epsilon^*$ RMSE= 0.0421
(c) TV_S, $\epsilon = 0.08\epsilon^*$ RMSE= 0.0447
(d) TV_S, $\epsilon = 0.09\epsilon^*$ RMSE= 0.0505
(e) TV_S, $\epsilon = 0.1\epsilon^*$ RMSE= 0.0564

(f) TNV, $\epsilon = 0.06\epsilon^*$ RMSE= 0.0410
(g) TNV, $\epsilon = 0.07\epsilon^*$ RMSE= 0.0398
(h) TNV, $\epsilon = 0.08\epsilon^*$ RMSE= 0.0392
(i) TNV, $\epsilon = 0.09\epsilon^*$ RMSE= 0.0392
(j) TNV, $\epsilon = 0.1\epsilon^*$ RMSE= 0.0405

Figure 8: Bone basis image, bone ROI comparison with channel-by-channel TV (top) and TNV (bottom). Grayscale window in $\mathrm{g/mL}^{-1}$ $[0.00, 0.52]$.

degree of noise suppression (high values of ϵ) without significantly deteriorating bone or soft-tissue structures. However, when the naïve channel-by-channel TV is used, these same ϵ values eventually lead to significant blurring artifacts. We also point out that even though the soft-tissue, 80 kVp, and 140 kVp images have edges that are not present in the bone density image, the TNV coupling does not falsely propagate these edges into the bone channel.

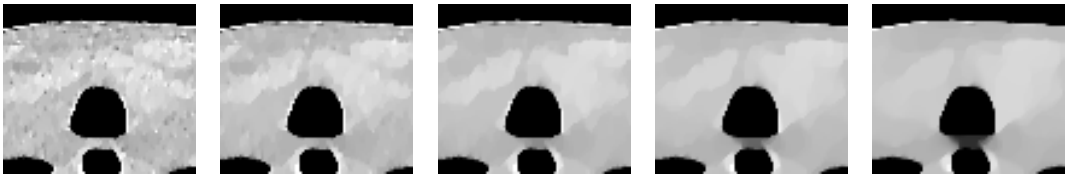

(a) TV_S, $\epsilon = 0.06\epsilon^*$ RMSE= 0.0445
(b) TV_S, $\epsilon = 0.07\epsilon^*$ RMSE= 0.0410
(c) TV_S, $\epsilon = 0.08\epsilon^*$ RMSE= 0.0383
(d) TV_S, $\epsilon = 0.09\epsilon^*$ RMSE= 0.0376
(e) TV_S, $\epsilon = 0.1\epsilon^*$ RMSE= 0.0400

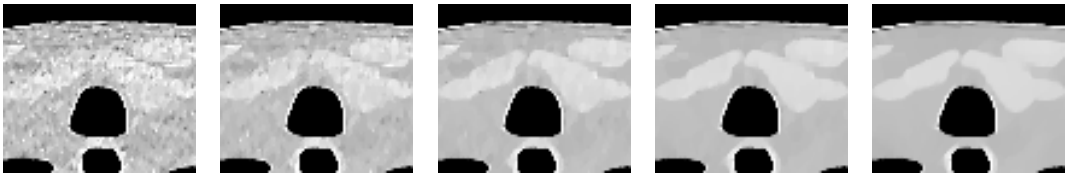

(f) TNV, $\epsilon = 0.06\epsilon^*$ RMSE= 0.0431
(g) TNV, $\epsilon = 0.07\epsilon^*$ RMSE= 0.0396
(h) TNV, $\epsilon = 0.08\epsilon^*$ RMSE= 0.0391
(i) TNV, $\epsilon = 0.09\epsilon^*$ RMSE= 0.0386
(j) TNV, $\epsilon = 0.1\epsilon^*$ RMSE= 0.0372

Figure 9: soft-tissue basis image, soft-tissue ROI comparison with channel-by-channel TV (top) and TNV (bottom). Grayscale window in $\mathrm{g/mL}^{-1}$ $[0.80, 1.07]$

5 Summary

We have described a framework for jointly reconstructing multi-channel spectral CT images using a generalized vectorial regularizer. Specifically, we presented a multi-channel generalization of the total variation which couples the different image channels by encouraging their gradient vectors to point in a parallel or anti-parallel direction. Preliminary results suggest that this coupling may allow for greater noise suppression with less blurring of image structures. This regularization strategy can be used to reconstruct raw energy bin data, synthetic basis-material images, or even both simultaneously, such as in our hybrid reconstruction study. This hybrid approach may allow for better noise suppression in the synthetic material images, which typically suffer from elevated noise levels. We believe that the proposed TNV provides an elegant way to regularize multi-channel images and may be a useful complement to previously published statistical, spectral CT reconstruction methods.

References

[1] D. S. Rigie and P. J. La Rivière, "A generalized vectorial total-variation for spectral CT reconstruction," Proceedings of the Third International Conference on Image Formation in X-ray Computed Tomography, vol. 1, no. 1, pp. 9–12, 2014.

[2] K. Holt, "Total Nuclear Variation and Jacobian Extensions of Total Variation for Vector Fields.," IEEE transactions on image processing : a publication of the IEEE Signal Processing Society, vol. 23, pp. 3975–3989, June 2014.

[3] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," Physica D: Nonlinear Phenomena, vol. 60, pp. 259–268, 1992.

[4] T. F. Chan and C. K. Wong, "Total variation blind deconvolution.," IEEE transactions on image processing : a publication of the IEEE Signal Processing Society, vol. 7, pp. 370–375, 1998.

[5] T. Komatsu and T. Saito, "Super-resolution sharpening-demosaicking method for removing image blurs caused by An optical low-pass filter," in Proceedings - International Conference on Image Processing, ICIP, vol. 1, pp. 845–848, 2005.

[6] E. J. Candès, J. Romberg, and T. Tao, "Robust Uncertainty Principles : Exact Signal Frequency Information," IEEE Transactions on Information Theory, vol. 52, no. 2, pp. 489–509, 2006.

[7] E. Sidky, C. Kao, and X. Pan, "Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT," Journal of X-ray Science and Technology, vol. 14, pp. 119–139, 2006.

[8] J. Bian, J. Wang, X. Han, E. Y. Sidky, L. Shao, and X. Pan, "Optimization-based image reconstruction from sparse-view data in offset-detector CBCT.," Physics in medicine and biology, vol. 58, pp. 205–30, 2013.

[9] J. Song, Q. H. Liu, G. A. Johnson, and C. T. Badea, "Sparseness prior based iterative image reconstruction for retrospectively gated cardiac micro-CT.," Medical physics, vol. 34, pp. 4476–4483, 2007.

[10] Z. Tian, X. Jia, and K. Yuan, "Low-dose CT reconstruction via edge-preserving total variation regularization," Physics in medicine and . . . , vol. 56, no. 18, pp. 5949–5967, 2011.

[11] G. Sapiro, "Color Snakes," Computer Vision and Image Understanding, vol. 68, pp. 247–253, 1997.

[12] P. Blomgren and T. F. Chan, "Color TV: Total variation methods for restoration of vector-valued images," IEEE Transactions on Image Processing, vol. 7, pp. 304–309, 1998.

[13] B. Goldluecke, E. Strekalovskiy, and D. Cremers, "The Natural Vectorial Total Variation Which Arises from Geometric Measure Theory," SIAM Journal on Imaging Sciences, vol. 5, pp. 537–563, 2012.

[14] S. Lefkimmiatis, A. Roussos, M. Unser, and P. Maragos, "Convex Generalizations of Total Variation Based on the Structure Tensor with Applications to Inverse Problems," in Scale Space and Variational Methods in Computer Vision 4th International Conference, SSVM 2013, Schloss Seggau, Leibnitz, Austria, June 2-6, 2013. Proceedings (A. Kujiper, K. Bredies, T. Pock, and H. Bischof, eds.), pp. 48–60, Springer Berlin Heidelberg, 4th ed., 2013.

[15] D. Keren and A. Gotlib, "Denoising Color Images Using Regularization and "Correlation Terms"," Journal of Visual Communication and Image Representation, vol. 9, no. 4, pp. 352–365, 1998.

[16] K. M. Holt, "Angular regularization of vector-valued signals," in ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings, pp. 1105–1108, 2011.

[17] J. A. Fessler, I. A. Elbakri, P. Sukovic, and N. H. Clinthorne, "Maximum-likelihood dual-energy tomographic image reconstruction," SPIE Medical Imaging, vol. 4684, pp. 38–49, 2002.

[18] Y. Long and J. A. Fessler, "Multi-Material Decomposition Using Statistical Image Reconstruction for Spectral CT.," IEEE transactions on medical imaging, vol. 33, pp. 1614–26, Aug. 2014.

[19] R. E. Alvarez and A. Macovski, "Energy-selective Reconstructions in X-ray Computerized Tomography," Phys Med Biol, vol. 21, pp. 733–744, 1976.

[20] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," Journal of Mathematical Imaging and Vision, vol. 40, pp. 120–145, 2011.

[21] W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. W. Tsui, "4D XCAT phantom for multimodality imaging research.," Medical physics, vol. 37, pp. 4902–4915, 2010.

[22] X. Han, J. Bian, E. L. Ritman, E. Y. Sidky, and X. Pan, "Optimization-based reconstruction of sparse images from few-view projections," Physics in Medicine and Biology, vol. 57, pp. 5245–5273, 2012.

[23] Q. Xu, A. Sawatzky, M. A. Anastasio, and C. O. Schirra, "Sparsity-regularized image reconstruction of decomposed K-edge data in spectral CT.," Physics in medicine and biology, vol. 59, pp. N65–79, May 2014.

[24] J. Cai, E. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," SIAM Journal on Optimization, pp. 1–28, 2010.

[25] E. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?," Journal of the ACM (JACM), 2011.

[26] E. Y. Sidky, J. H. Jørgensen, and X. Pan, "Convex optimization problem prototyping for image reconstruction in computed tomography with the ChambollePock algorithm," Physics in Medicine and Biology, vol. 57, pp. 3065–3091, 2012.

[27] N. Parikh and S. Boyd, "Proximal algorithms," Foundations and Trends in Optimization, vol. 1, no. 3, pp. 123–231, 2013.

[28] C. A. Bouman and K. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," IEEE Transactions on Image Processing, vol. 5, pp. 480–492, 1996.

[29] E. Roessl and C. Herrmann, "Cramér-Rao lower bound of basis image noise in multiple-energy x-ray imaging.," Physics in medicine and biology, vol. 54, pp. 1307–18, Mar. 2009.

[30] C. O. Schirra, E. Roessl, T. Koehler, B. Brendel, A. Thran, D. Pan, and R. Anastasio, Mark A.and Proksa, "Statistical reconstruction of material decomposed data in spectral CT.," IEEE transactions on medical imaging, vol. 32, pp. 1249–57, July 2013.

[31] A. Sawatzky, Q. Xu, C. O. Schirra, and M. A. Anastasio, "Proximal ADMM for Multi-Channel Image Reconstruction in Spectral X-ray CT.," IEEE transactions on medical imaging, vol. 33, pp. 1657–1668, Aug. 2014.

Appendices

A Deriving the saddle-point problem

To derive the saddle-point formulation of our optimization problem, we used two fundamental results of convex analysis to "dualize" the data fidelity constraint as well as the VTV term. The transformation of the VTV term follows straightforwardly from the definition of the so-called "dual-norm," which is defined by

$$\|x\|' = \sup_{z} \ \langle x, z\rangle \ \text{ s.t. } \|z\| \leq 1. \tag{25}$$

Every norm $\|.\|$ has an associated dual-norm $\|.\|'$, that obeys this relationship. The ℓ_2 norm utilized in the scalar TV is self-dual, while the nuclear norm and spectral norm form a dual pair.

Dualizing the channel-by-channel TV Recall that the channel-by-channel TV can be written as

$$\mathrm{TV}_S(\boldsymbol{u}) = \sum_{m,\ell} \left\| [D\boldsymbol{u}_m]_\ell \right\|_2 . \tag{26}$$

The definition of the dual-norm allows us to rewrite this as

$$\begin{aligned} \mathrm{TV}_S(\boldsymbol{u}) &= \sum_{m,\ell} \sup_{\boldsymbol{z}} \ \langle [D\boldsymbol{u}_m]_\ell , [\boldsymbol{z}_m]_\ell \rangle \ \text{ s.t. } \| [\boldsymbol{z}_m]_\ell \|_2 \leq 1 \\ &= \sup_{\boldsymbol{z}} \ \langle D\boldsymbol{u}, \boldsymbol{z} \rangle \ \text{ s.t. } \ \left\| [\boldsymbol{z}_m]_\ell \right\|_2 \leq 1 \ \ \forall m \in 1, \ldots, \mathcal{M} \ \ ; \ell \in 1, \ldots, \mathcal{L} \end{aligned} \tag{27}$$

which is equivalent to (18).

Dualizing the proposed TNV The proposed TNV is

$$\mathrm{TNV}(\boldsymbol{u}) = \sum_{\ell=1}^{\mathcal{L}} \| \, [D\boldsymbol{u}]_\ell \, \|_\star . \tag{28}$$

Substituting the definition of the dual-norm, we get

$$\begin{aligned} \mathrm{TNV}(\boldsymbol{u}) &= \sum_{\ell} \sup_{\boldsymbol{z}} \ \langle [D\boldsymbol{u}]_\ell , [\boldsymbol{z}]_\ell \rangle \ \text{ s.t. } \sigma_{\max} \left([\boldsymbol{z}]_\ell \right) \leq 1 \\ &= \sup_{\boldsymbol{z}} \ \langle D\boldsymbol{u}, \boldsymbol{z} \rangle \ \text{ s.t. } \sigma_{\max} \left([\boldsymbol{z}]_\ell \right) \leq 1 \ \ \forall \ell \in 1, \ldots, \mathcal{L}, \end{aligned} \tag{29}$$

which is equivalent to (20). We have used the fact that the spectral norm (maximum singular value) is dual to the nuclear norm.

Dualizing the data-fidelity term using the Fenchel conjugate To dualize the data-fidelity term, we make use of the Fenchel conjugate. For a convex, lower semi-continuous function f, the conjugate f^* is defined by

$$f^*(x) = \sup_{x'} \ \langle x, x' \rangle - f(x'), \tag{30}$$

where it is also true that $(f^*)^* = f$. Using this definition, it is easy to show that the following functions form a conjugate pair,

$$\delta_\epsilon(x) \Leftrightarrow \epsilon \|W^{-1/2}x\|_2 \tag{31}$$

where $\delta_\epsilon(x)$ is the indicator function, defined by (16). This along with the definition of the Fenchel conjugate leads directly to Eqn. 17.

B The proximal map of $\epsilon\sigma\|W^{-1/2}\boldsymbol{q}\|_2$

The update equations for Algorithm 2 involve evaluating the proximal map of $\epsilon\sigma F_q \equiv \epsilon\sigma\|W^{-1/2}\boldsymbol{q}\|_2$, which does not admit a closed form. However, we will now show how it can still be efficiently evaluated using a simple root-finding procedure. First we explicitly write out the proximal mapping as

$$\mathbf{prox}_{\epsilon\sigma F_q}(\boldsymbol{q}') = \underset{\boldsymbol{q}}{\operatorname{argmin}} \ \|W^{-1/2}\boldsymbol{q}\|_2 + \frac{1}{2\epsilon\sigma}\|\boldsymbol{q} - \boldsymbol{q}'\|_2^2. \tag{32}$$

Next, we compare this optimization problem to a slightly easier problem,

$$\underset{\boldsymbol{q}}{\operatorname{argmin}} \ \frac{1}{2}\|W^{-1/2}\boldsymbol{q}\|_2^2 + \frac{1}{2\lambda\sigma}\|\boldsymbol{q} - \boldsymbol{q}'\|_2^2, \tag{33}$$

and note that for some choice of ϵ and λ, these objectives both have the same set of level curves. We can find the solution to this problem by setting the gradient equal to zero, and for a symmetric weighting matrix W, it is given by

$$\boldsymbol{q}^* = \left(I + \sigma\lambda W^{-1}\right)^{-1}\boldsymbol{q}'. \tag{34}$$

If the gradient of this easier problem is equal to that of the original problem, then $\boldsymbol{q}^*$ must be an optimizer of the original problem as well. Momentarily, we ignore the non-differentiable point $\boldsymbol{q} = \mathbf{0}$, and set the gradients equal, which yields

$$\begin{aligned} \epsilon &= \lambda\|W^{-1/2}\boldsymbol{q}^*\|_2. \\ &= \lambda\|W^{-1/2}\left(I + \sigma\lambda W^{-1}\right)^{-1}\boldsymbol{q}'\|_2. \end{aligned} \tag{35}$$

In this work, we will only consider the case where W is a diagonal matrix, which allows us to simplify this expression.

$$\epsilon^2 - \lambda^2 \sum_{j=1}^{J} \frac{W_{jj}}{\left(W_{jj} + \sigma\lambda\right)^2}[\boldsymbol{q}']_j^2 = 0. \tag{36}$$

The index j is a linear index over every element of $\boldsymbol{q}$. Now, we just need to perform a 1D search for a positive value of λ that satisfies this equation, which can be done very efficiently using Newton's method or a variety of other algorithms. Once this root, λ^*, is found, we simply plug it back into (35):

$$\mathbf{prox}_{\epsilon\sigma F_q}(\boldsymbol{q}') = \begin{cases} \left(I + \sigma\lambda^* W^{-1}\right)^{-1}\boldsymbol{q}' & \|W^{1/2}\boldsymbol{q}'\|_2 > \sigma\epsilon \\ 0 & \|W^{1/2}\boldsymbol{q}'\|_2 \leq \sigma\epsilon \end{cases}. \tag{37}$$

Note that because of the non-differentiability of $\|W^{-1/2}\boldsymbol{q}\|_2$ at $\boldsymbol{q} = 0$, there will not always be a solution to equation 36. Specifically, if $\sigma\epsilon > \|W^{1/2}\boldsymbol{q}'\|_2$, then there is no viable root. In this case, we need not perform

the root-finding procedure, because the optimum must occur at $\boldsymbol{q} = \mathbf{0}$.

C Implementation of the projection operators $\Pi_{\mathcal{S}/N}$

$\boldsymbol{\Pi_{\mathcal{S}}}$ Now, we will describe the Euclidean projection onto the set $\mathcal{S}$ defined in (19). Consider projecting $\boldsymbol{z} \in \mathbb{R}^{(\mathcal{M}\times 2)\cdot\mathcal{L}}$ onto $\mathcal{S}$. We can define this operation element-wise on each vector $[\boldsymbol{z}_m]_\ell \in \mathbb{R}^2$ corresponding to the m^{th} spectral channel and the ℓ^{th} pixel-location. The projection of this element is given by

$$\Pi_{\mathcal{S}}\left([\boldsymbol{z}_m]_\ell\right) = \begin{cases} [\boldsymbol{z}_m]_\ell & \|[\boldsymbol{z}_m]_\ell\|_2 \leq 0 \\ \frac{[\boldsymbol{z}_m]_\ell}{\|[\boldsymbol{z}_m]_\ell\|_2} & \text{otherwise} \end{cases}. \tag{38}$$

For every image channel m and pixel location ℓ, we simply project the vector $[\boldsymbol{z}_m]_\ell$ onto the unit ball.

$\boldsymbol{\Pi_N}$ The projection onto set N defined by (21) is slightly more complicated. We define this operation element-wise on each $\mathcal{M} \times 2$ matrix $[\boldsymbol{z}]_\ell$. This time we need to threshold the eigenvalues of $[\boldsymbol{z}]_\ell$. This projection can be succinctly described by

$$\Pi_N\left([\boldsymbol{z}]_\ell\right) = U\Sigma_p V^T, \tag{39}$$

where $U\Sigma V^T$ is the SVD of $[\boldsymbol{z}]_\ell$, and

$$\Sigma_p = \operatorname{sgn}(\Sigma)\min\left(|1|, |\Sigma|\right) \tag{40}$$

is the thresholded version of Σ. To form Σ_p the singular values on the diagonal of Σ are simply thresholded so that their magnitude does not exceed 1. An equivalent projection formula that leads to a much more computationally efficient solution is given by

$$\Pi_N\left([\boldsymbol{z}]_\ell\right) = [\boldsymbol{z}]_\ell V\Sigma^\dagger\Sigma_p V^T. \tag{41}$$

The quantity $\Sigma^\dagger$ is the pseudo-inverse of Σ. Since we are only working with 2 spatial dimensions the matrix V will by a 2×2 square matrix. Therefore, to compute the projection according to (41) we only need to compute the eigenvalues and eigenvectors of a 2×2 matrix, for which a very simple closed form is available. This can also be done very efficiently for 3D images, where V will be 3×3. Because of this, the update equations that result from the proposed vectorial TV are only trivially more expensive than the channel-by-channel TV, and the projection/backprojection operations are likely to swamp this difference.